

NAS Grid Benchmarks:

A Tool for Measuring Performance of Computational Grids

Michael Frumkin (NAS Division)
 Rob F. Van der Wijngaart (Computer Sciences Corp.)
 NASA Ames Research Center
 Moffett Field, CA

GGF3, Oct 7-10, 2001

*Philip Dormer Stanhope, Earl of Chesterfield (1746):
 A job worth doing is worth doing well*

*Nerd's addendum:
 A distributed job worth doing is worth doing well*

*Nerd's addendum's addendum:
 A distributed job worth doing is worth measuring*

Frumkin, Van der Wijngaart

GGF3, Oct 7-10, 2001

Contents

1. Brief history of benchmarking at NAS
2. Measuring Grid performance
3. Preliminary results

Frumkin, Van der Wijngaart

GGF3, Oct 7-10, 2001

Brief history of benchmarking at NAS

NAS Parallel Benchmarks (NPB)

Goal: MMM Measure efficiency of modern (parallel) machines running scientific application

Audience: M User: application engineer, researcher
 Buyer + vendor

Contents: 5 kernels, 3 pseudo apps (implicit CFD)

Approach: NPB1: Paper-and-pencil specs – '91
 NPB2: Source code implementations (F77/C/MPI) – '97
 PBN: Source code implementations (HPF/OpenMP/Java) – '99

Frumkin, Van der Wijngaart

GGF3, Oct 7-10, 2001

NAS Parallel Benchmark codes

Kernels: M MMM

- EP Random-number generator
- IS Integer sort
- CG Conjugate gradient
- MG Multigrid method for Poisson eqn
- FT Spectral method (FFT) for Laplace eqn

Pseudo apps:

- BT ADI; Block-Tridiagonal systems
- SP ADI; Scalar Pentadiagonal systems
- LU Lower-Upper symmetric Gauss-Seidel

Frumkin, Van der Wijngaart

GGF3, Oct 7-10, 2001

Measuring Grid performance

NAS Grid Benchmarks (NGB)

Goals: MM Measure efficiency + functionality of modern Grids running (distributed) scientific application (includes public network performance)

Audience: Grid user: application engineer, researcher
 Audience: Grid software/infrastructure developer

Contents: 4 compound tasks: 3 pseudo apps, 2 kernels, all from NPB

Approach: NGB1: Paper-and-pencil specs – '01
 NGB2: Source code implementations (NPB2/PBN + Globus/Legion/CORBA/ Grid Engine/Condor/Java ...)

Frumkin, Van der Wijngaart

GGF3, Oct 7-10, 2001

Grid benchmark requirements

- Tests computational aspects of environment
- Is representative of scientific computing tasks
- Uses basic Grid services
- Is not intrusive (no throughput stress testing)
- Contains communicating processes
- Does significant communication
- Is verifiable (deterministic, not interactively steered)
- Needs no initialization data files
- Is fair (no favorite before dust settles)

7

GGF3, Oct 7-10, 2001

Grid benchmark construction

- Provide synthetic Grid app. for scientific computing:
 - Data Flow Graph (DFG) coupling scientific codes
- Specify:
 - abstract services: authenticate, create task, communicate
 - problem sizes (classes): S, A, B, C, ...
- Do not specify:
 - Mapping, scheduling, fault tolerance, data security
- Measure and report turnaround time

8

GGF3, Oct 7-10, 2001

Nodes of NGB Data Flow Graphs

Modified NPB mesh-based codes SP, BT, LU, MG, FT:

- Well-accepted, well-studied, widely portable
- Solid verification procedures exist
- Parallel versions available: MPI/OpenMP/HPF/Java
- No data files required, but ...
- ... output of one NPB may be input for another
- Symbolize components of scientific apps:
 - BT, SP, LU: scientific comps (flow solvers)
 - MG: post-processing (data smoother)
 - FT: visualization (spectral analysis)

9

GGF3, Oct 7-10, 2001

NGB Data Flow Graphs, Sample size

Embarrassingly Distributed (ED)

Parameter study

Helical Chain (HC)

Cyclic process (restart)

10

GGF3, Oct 7-10, 2001

NGB Data Flow Graphs, Sample size

Visualization Pipe (VP)

Visualization cycle

Mixed Bag (MB)

Unbalanced chain

11

GGF3, Oct 7-10, 2001

Approximate rationale for NGB DFG classes

- S: for testing only; completes in seconds
- A: reasonable size eng. problem; fits on modern SSI machine without time sharing processes
- B: challenging problem size for advanced SSI machine
- C: won't fit on most advanced SSI machines

Considerations:

- Keep # verification values limited
- Provide significant range of problem sizes
- Complete in reasonable time (s-hrs on critical path)
- Benchmarks non-converging, numerically stable
- Ensure $T(\text{compute}) \sim T(\text{communicate})$

12

GGF3, Oct 7-10, 2001

Larger NGB DFG classes

$(ED.X \mid X \in \{A,B,C\}) = \Sigma \{16,64,256\} SP.X$
 $(HC.X \mid X \in \{A,B,C\}) = \Pi \{3,6,12\} (BT.X, SP.X, LU.X)$
 $(VP.X \mid X \in \{A,B,C\}) = \Pi \{9,18,36\} (BT.X, MG.X, FT.X)$

VP

MB


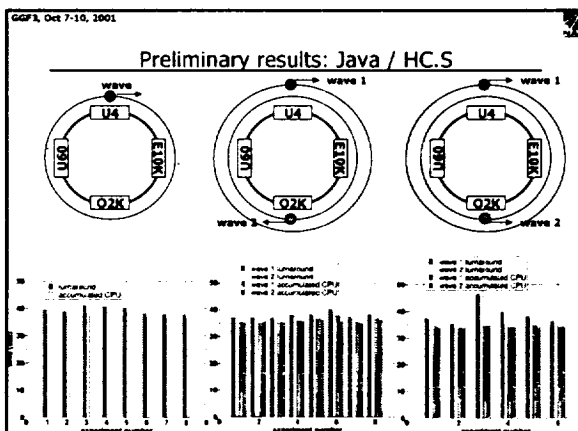
13

GGF3, Oct 7-10, 2001

Preliminary results

- Pilot implementations in Java and Globus under development.
- Preliminary results for ED.S and HC.A on proto-cluster and Origin cluster, resp.

14



GGF3, Oct 7-10, 2001

NGB into the future

Possible use of NGBs; interpretation of results

Primitive Grid services

- Functionality
- Consistency among runs

Complete resources reservation

- Variation of single resource

Definition of uniform Grid currency (G\$)

- Cost of NGB performance per G\$

Grid maturity increases

NUG complexity: $O(n)$ CPU(27)-3,500, CPU(28)-20,000, CPU(30)-100,000
 Wild guess: NUG complexity: $n/4 \cdot n$ $Q(30)/Q(28) \sim 60$

17

GGF3, Oct 7-10, 2001

Errata

- No NGB class W
- Only verification in Report node
- LP \rightarrow HC, CP \rightarrow VP

18